ABSTRACT
         A prescriptive model of problem-solving performance
is described. The aim of the model is to specify explicitly how
problems should be described initially in such a form as to
facilitate their subsequent solution. According to the model, the
process of initial problem description can be decomposed into two
successive stages. The first uses general domain-independent
knowledge to generate a "basic description" of a problem, the purpose
of which is to identify explicitly the information specified and
wanted in the problem, to introduce useful symbols, and to express
the relevant information in various convenient symbolic
representations. This basic description is then used to generate a
"theoretical description" of the problem, that is, a deliberate
redescription of the problem in terms of special concepts provided by
the domain-specific knowledge base, involving the particular entities
of interest in the problem, describing these entities in terms of
concepts specifically provided by the knowledge base, and exploiting
the known properties of these descriptive concepts. Focusing on
mechanics (physics), the knowledge needed to generate theoretical
problem descriptions is discussed, followed by experimental methods
for testing the prescriptive model and selected results obtained from
such experiments. (Author/JN)

Paper to be presented at a session on "Acquisition of Skills in Solving Physics Problems: Strategies and Representations" at the annual meeting of the American Educational Research Association in New York City, 23 March 1982.

COGNITIVE MECHANISMS FACILITATING HUMAN PROBLEM SOLVING IN PHYSICS:

FORMULATION AND ASSESSMENT OF A PRESCRIPTIVE MODEL*

F. Reif and Joan I. Heller

University of California, Berkeley

BASIC GOALS AND APPROACH

Problem solving is a centrally important activity in any science. In order to teach such problem solving systematically, it is first necessary to understand how good problem solving performance is achieved. Such an understanding, together with an understanding of the performance of novice students before any instruction, allows one then to undertake systematic efforts to design explicit instructional methods for teaching problem solving.

In trying to understand cognitive mechanisms leading to good problem solving, it is instructive to study what expert problem solvers actually do (Larkin & Reif, 1979; Larkin, McDermott, Simon, & Simon, 1980; Chi, Feltovich, and Glaser, 1981), but unwise to restrict one's focus in this way. In particular, experts do not necessarily always perform optimally. Furthermore, the performance to be achieved by students as a result of instruction can not merely mimic expert performance. Indeed, students must often use explicit procedures to achieve performance which experts accomplish almost automatically by recognizing patterns with which they have become familiar as a result of years of experience.

2

Accordingly, the aim of our work has been to study effective human per-
formance from a more general "prescriptive" point of view which transcends
the description of naturalistically occurring phenomena. In particular, our
aim has been to specify explicitly the kinds of underlying knowledge leading
to good human problem solving performance in a realistic scientific domain,
without necessarily trying to simulate what actual experts do. Such a pre-
scriptive point of view is clearly more general than a descriptive one. For
example, although a prescriptive theoretical model of good performance may
be partly suggested by naturalistic observations of expert behavior, it may
also be suggested by purely theoretical task analyses. Correspondingly, the
criterion of validity of a prescriptive model of good performance is solely
that it lead to predictably effective performance when implemented by a human
subject, even if it does not simulate closely what actual experts do.

An analogy might help to clarify the distinction between a more general
prescriptive point of view and a descriptive one. A hypothetical cognitive
scientist, working in the year 100 AD to formulate a model of good perfor-
mance in arithmetical problem solving, might have suggested the use of the
modern place-value representation of numbers. The resulting model would have
led to good arithmetical performance. On the other hand, it would have been
a very poor descriptive model of the performance of experts, all of whom used
Roman numerals at that time.

A prescriptive approach is more general than a naturalistic one because
it allows greater room for human invention and experimental manipulation. It
can also be very useful for identifying essential knowledge required for good
performance and can thereby help to explicate expert knowledge which is often
largely tacit. From an applied point of view, a prescriptive approach is

also centrally important for instruction and for attempts to improve human performance.

In trying to specify explicitly the kinds of knowledge leading to good performance in a scientific domain, we have focused our attention on good problem solving performance in basic college-level physics, specifically in the field of mechanics. Such scientific problem solving is realistically complex, yet sufficiently simple to be amenable to systematic investigation. Furthermore we have presupposed that the human subjects, who engage in such problem-solving tasks, possess well-developed human capabilities such as a knowledge of English, a knowledge of simple algebra, and a knowledge of basic physics principles. On the other hand, we do not assume that such human subjects possess more sophisticated or strategic forms of knowledge required for problem solving. Indeed, our main aim is precisely to specify explicitly such strategic procedures and forms of knowledge organization leading to good problem solving performance.

INITIAL PROBLEM DESCRIPTION

Our prescriptive model (Reif & Heller, 1981), which aims to specify these more sophisticated forms of knowlege leading to good problem solving, involves knowledge of different levels of generality, i.e., domain-specific knowledge (e.g., specific knowledge about mechanics) embedded in more general domain-independent knowledge. The domain-independent knowledge includes general problem-solving procedures subdividing the problem-solving process into successive stages. These stages include (a) the generation of an initial problem description designed to facilitate subsequent construction of the problem solution; (b) the actual construction of the solution, including procedures for making judicious decisions facilitating search; and (c)

4

the subsequent assessment and improvement of the solution. The preceding general procedures are to be used in conjunction with a "knowledge base" containing specific knowledge about the particular domain of interest. This knowledge base should have general characteristics which facilitate the implementation of the general problem-solving procedures, e.g., it should contain certain kinds of knowledge and be organized in certain ways. Such important characteristics of the knowledge base must also be specified by a prescriptive model of good performance.

The description of a problem is of central importance since the solution of a problem can be implemented only if it has been formulated in an appropriate representation. Indeed, the <u>initial</u> description of a problem can crucially determine how easily a problem can subsequently be solved, or whether it can be solved at all. Actual observations of experts provide, however, rather little direct information about the description process, since experts tend to describe problems rapidly and seemingly automatically on the basis of large amounts of tacit knowledge.

Our prescriptive model of good problem-solving performance aims to specify explicitly how problems should be described initially in such a form as to facilitate their subsequent solution. According to this model, the process of initial problem description can be decomposed into two successive stages. The first of these uses general domain-independent knowledge to generate a "basic description" of a problem. The purpose of such a basic description is merely to identify explicitly the information specified and wanted in the problem, to introduce useful symbols, and to express the relevant information in various convenient symbolic representations (e.g., pictorial as well as verbal forms).

5

This basic description is then used to generate a "theoretical description" of the problem, i.e., a deliberate redescription of the problem in terms of special concepts provided by the domain-specific knowledge base. In particular, the generation of such a theoretical problem description involves identifying the particular entitites of interest in the problem; describing these entitites in terms of concepts specifically provided by the knowledge base; and exploiting the known properties of these descriptive concepts. Such a deliberate redescription of a problem in terms of the concepts provided by the knowledge base greatly facilitates the subsequent search for a solution since all principles in the knowledge base are expressed in terms of these particular concepts and become then readily accessible.

The preceding comments imply that a well-structured knowledge base about any domain should have characteristics which facilitate the useful description of any situation within this domain. In particular, an effective knowledge base should specify the particular entities of interest in the specified domain; the concepts most useful for describing these entities; and the properties of these concepts. Furthermore, the knowledge base should contain explicit guidelines specifying procedures for using the preceding knowledge to describe any situation within the specified domain.

Because of limitations of time and space, the present paper focuses attention primarily on explicit procedures for generating effective theoretical problem descriptions. Needless to say, this is only a small, but important, part of our more-encompassing model of good problem-solving performance. Restricting our attention to the particular domain of mechanics, we shall first specify, from a prescriptive point of view, the knowledge needed to generate good theoretical problem descriptions in this domain. Then we shall

.discuss experimental methods for testing such a theoretical model and some

of the results obtained by such experiments.

THEORETICAL PROBLEM DESCRIPTION IN MECHANICS

The preceding general remarks about theoretical problem description can

now be exemplified in the particular science of mechanics by identifying the

particular entities of interest in this domain, the special concepts used to

describe these entities, and the particular properties of these concepts.

The knowledge base of mechanics specifies that the entities of interest

in this domain are particles and systems consisting of several particles,

(e.g., strings, rigid bodies,...). As indicated in Figure 1, the knowledge

base introduces two different classes of special concepts to describe such

particles, namely "intrinsic descriptors" and "interaction descriptors". The

intrinsic descriptors describe individual particles. Some of these descrip-

tors (such as "mass") merely characterize any particle; the others are

"motion descriptors" (such as "position", "velocity", "acceleration") which

are used to describe the motion of any particle. By contrast, the inter-

action descriptors do not describe individual particles, but the interaction

between such particles. For example, the "force" exerted on a particle by

some other particle is one such interaction descriptor, "potential energy" is

another one.

Insert Figure 1 about here

The knowledge base for mechanics specifies important properties of the

preceding descriptors. In particular, "interaction laws" specify how the

interaction descriptors are related to the intrinsic descriptors of the

interacting particles (e.g., how the force on one particle by another is

KNOWLEDGE BASE FOR MECHANICS

ENTITIES: particles and systems thereof

INTRINSIC DESCRIPTORS	INTERACTION DESCRIPTORS

characteristics ⋀ motion
 	
mass   position
         velocity
         acceleration
	force, potential energy,...

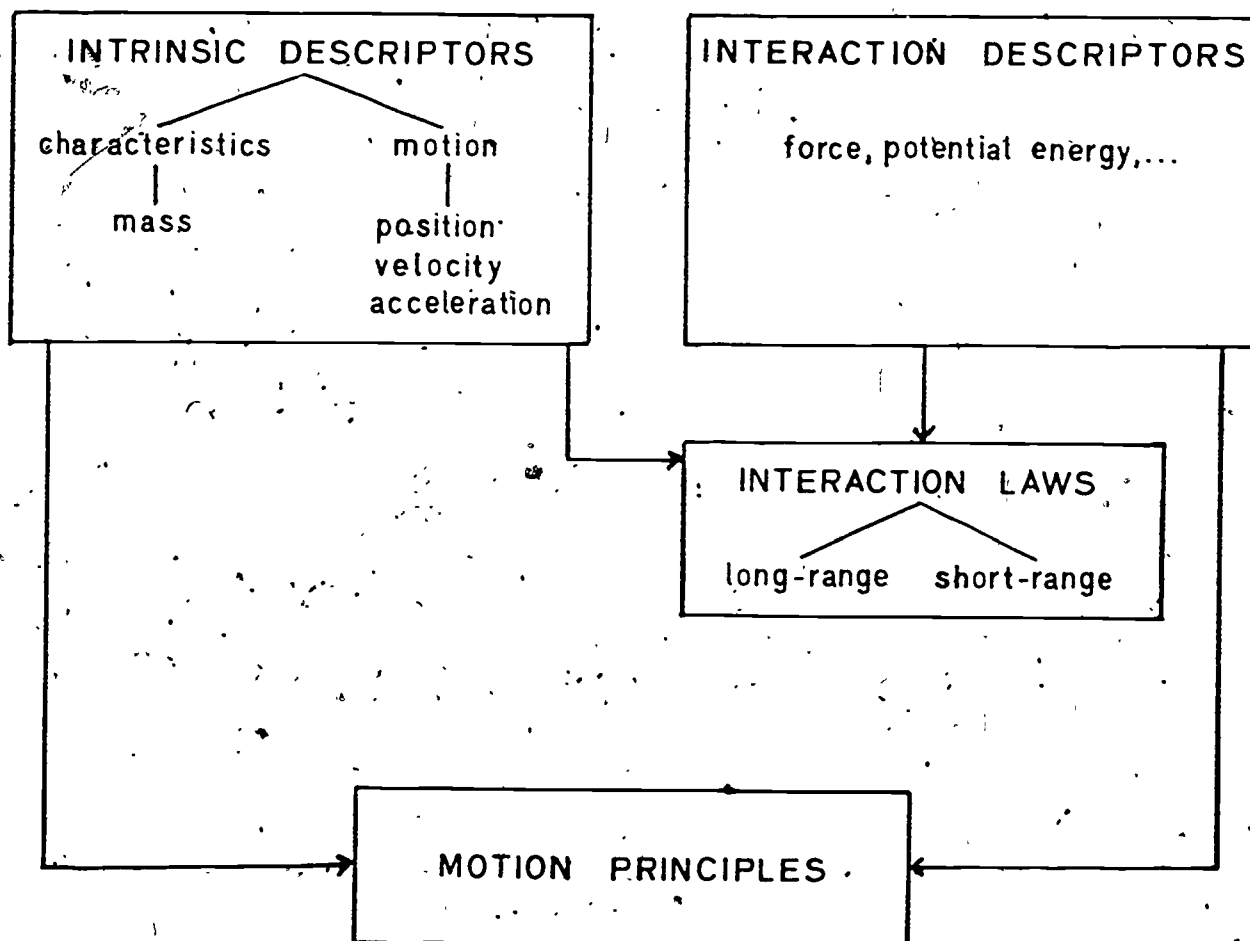INTERACTION LAWS

long-range   short-range

MOTION PRINCIPLES

Figure 1

8

related to the characteristics of the particles and to their relative positions). Such interaction laws are specified for various kinds of interactions encountered in nature, but can be classified into the following two types of interactions: Some of these interactions are "long-range" because they are appreciable even if the interacting particles are separated by an appreciable distance. (The prime example is the gravitational interaction of a particle with the earth.) The other interactions are "short-range" because they are only appreciable when the interacting particles are so close that they "touch" each other. (Examples are the interaction of a particle in contact with a string or with the surface of a solid object.)

Lastly, the knowledge base for mechanics specifies important "motion principles" which specify how the motion descriptors of particles change with time as a result of the interaction between particles (e.g., how the acceleration of a particle depends on the force on this particle by other particles). These motion principles provide the science of mechanics with its great predictive power.

The preceding factual knowledge in the knowledge base for mechanics can be accompanied by explicit rules specifying how this knowledge is to be used for generating an explicit theoretical description of any problem in mechanics. The main steps of this description procedure are the following:

(1) Identify the particles of interest at each time of interest.

(2) Describe the motion of each such particle by drawing a diagram indicating all available knowledge about the position, velocity, and acceleration of this particle.

(3) Describe the interaction of this particle by drawing a diagram indicating all forces on this particle. Do this as follows: (a) Identify all objects interacting with the given particle by long-range forces. (Ordinarily

9

this is just the earth interacting with the given particle by gravitational forces.) Indicate the corresponding forces and their properties, as specified by the interaction laws for long-range forces. (b) Identify all objects which <u>touch</u> the given particle. For each such touching object indicate the corresponding short-range forces and their properties, as specified by the interaction laws for short-range forces.

(4) Check that the description of motion and of forces is qualitatively consistent with known motion principles (e.g., that the acceleration of the particle has the same direction as the total force on it).

The preceding procedure, based on the knowledge base of mechanics, constitutes a prescriptive model of how to describe effectively any problem in mechanics. One would predict that the implementation of this description procedure by a human subject should lead to the following important consequences:

(1) It should lead to a very explicit initial description of any mechanics problem in terms of the special concepts of mechanics. In particular, this description should be appreciably more explicit than that apparent from observations of experts or than that presented in textbooks.

(2) The description procedure should help students avoid many common errors. For example, it should help avoid the common mistakes of omitting certain forces or of enumerating non-existent forces.

(3) The description procedure should sometimes lead to easier reformulations of problems. (For example, a question asking "when a string becomes slack" is automatically transformed into a question asking "when the force by the string becomes zero", a question which is much more easily interpreted and answered.)

(4) The explicit problem description generated by this procedure should
appreciably facilitate the subsequent solution of mechanics problems. In-
deed, the initial problem description can be the major difficulty in certain
problems and, once implemented, can make their subsequent solution fairly
trivial.

EXPERIMENTAL METHODS FOR TESTING A PRESCRIPTIVE MODEL

In the preceding sections we have outlined a prescriptive model whereby
a human subject should reliably be able to generate useful initial problem
descriptions. As mentioned previously, the ultimate criterion of validity
of such a model is not whether it simulates closely what actual experts do,
but whether it leads to predictably good performance. Accordingly, the
basic paradigm for testing the validity of such a prescriptive model is the
following: Induce a human subject to act in accordance with the prescriptive
model and observe whether the resulting performance has the predicted and
desired characteristics.

A particular way of implementing this general paradigm is in experiments
in which a human subject is induced to act under "external control". This
experimental procedure may be clarified by a familiar analogy, the situation
where a pilot lands his plane in bad weather while following directions from
an air-traffic controller on the ground. Under these conditions, a human
information processor (the pilot) makes extensive use of his sophisticated
knowledge, but relegates higher-level control of this knowledge to external
directions. This situation can be viewed as an experiment with the following
interesting characteristics: (1) It allows a separation of high-level con-
trol knowledge from lower-level implementation knowledge. For example, if
the plane were to crash, the information retrievable from the taped

conversations between pilot and ground control would allow one to distinguish whether the crash occurred as a result of sound control directions improperly executed by the pilot, or whether it occurred as a result of faulty control directions. By contrast, if a pilot crashed his plane while flying entirely under his own control, one could not distinguish whether the fault was in the pilot's higher-level control knowledge or lower-level implementation knowledge. (2) A set of control directions specifying how to land a plane can be viewed as a cognitive theory specifying how a human subject, with sophisticated human capabilities, can land a plane. In other words, such control directions would constitute a good theory of plane landing if, and only if, the correct execution of these directions leads to reliably effective landing of planes. (3) Such a validated theory of plane landing could ultimately be used as the basis of a theory of instruction for landing planes. Such an instructional theory would need to teach human subjects to internalize, and carry out independently, the control directions which had previously been external.

Let us now turn from our analogy to external-control experiments designed to test other prescriptive models of human performance, e.g., models of effective problem description. To carry out such experiments, one needs first to design a program consisting of step-by-step directions, and associated knowledge, whereby a human subject can be guided to act in accordance with a specified model of performance. For example, such a program might guide a human subject to execute the explicit description procedure discussed earlier. Such a program should be problem-independent, i.e., equally applicable to any problem in the specified domain. Furthermore, one must make sure that the directions in such a program are properly matched to the characteristics and pre-existing knowledge of the human subjects for whom the

program is designed. In particular, the individual directions specified by the program must be reliably interpretable and executable by the human subject. They must also be formulated at an appropriate level of detail, i.e., detailed enough to provide adequate guidance, but not so excessively detailed as to be burdensome or distracting to the subject.

In the actual experimental procedure an individual human subject is then asked to carry out specified tasks (e.g., the description and subsequent solution of various problems) by executing successively verbally stated directions according to a program specified by the model.. While so doing, the subject is asked to talk out loud about his or her thought processes and the whole session is tape-recorded. Detailed data can thus be gathered about the subject's written output and verbalized thought processes while responding to the external control directions.

Such detailed observations allow one to obtain the following kinds of information to test the proposed model of performance:

(1) One can ascertain whether the proposed model of good performance is, in fact, sufficient to lead to good performance. This can be done by determining whether subjects, working under external control in accordance with the model, do indeed achieve good performance. (Note that such experiments do not imply that the proposed model is unique since other models might also be sufficient, or even superior, in producing good performance.)

(2) One can verify that the prerequisite basic knowledge, which the model presupposes of human subjects, is by itself not sufficient to produce good performance. This can be done by letting subjects, with such knowledge, work without external guidance of the model and observing that the resulting performance is poor.

(3) One can ascertain whether selected features of the proposed model are, in fact, <u>necessary</u> to achieve good performance. This can be done by comparative experiments where human subjects work under external control of a modified model which lacks selected features of the proposed model of good performance. The predicted performance deficiencies should then occur.

(4) Finally, one can test whether the proposed model of good performance, when implemented, leads to specific predicted features in the resulting performance. For example, one can ascertain whether, and how, the occurrence of specific errors is prevented when human subjects act in accordance with the model.

It should be emphasized that the aim of such external-control experiments is to ascertain the merits of a proposed model of good performance, but <u>not</u> to teach. Subjects may, of course, learn incidentally while working under conditions of external control. However, such learning need not occur because external control directions do not become internalized. A subject, performing very well while working under external control, might thus revert to poor performance when external control knowledge is subsequently removed.

The following paper, by Heller and Reif, discusses in much greater detail the implementation and results of such external-control experiments designed to test a model for generating effective initial problem descriptions.

REFERENCES

Chi, M.T.H., Feltovich, P. J. & Glaser, R., Categorization and represention of physics problems by experts and novices. <u>Cognitive Science</u>, 1981, <u>5</u>, 121-152.

Larkin, J. H. & Reif, F., Understanding and teaching problem solving in physics. <u>European Journal of Science Education</u>, 1979, <u>1</u>, 191-203.

Larkin, J., McDermott, J., Simon, D. P., & Simon, H. A., Expert and novices

    performance in solving physics problems. Science, 1980, 208, 1335-1342.

Reif, F. & Heller, J. I., Knowledge structure and problem solving in physics.

    Report ES-12, Physics Department, University of California, 1981.